(1) To add a more formal probabilistic model than is used in the INTERNIST system which has greatly influenced our work. This will allow assumptions made by NESTOR to be explicit.

(2) To represent temporal knowledge explicitly and to use it in reasoning.

(3) To represent and use causal information to express the pathophysiology of diseases.

(4) To incorporate abstract diseases into the representation (e.g., it will be possible to include CANCER as a component of a user's hypothesis if this seems more appropriate at the time than specifying BRONCHOGENIC-CARCINOMA or MYELOMA or RENAL-CELL-CARCINOMA).

We have already developed several techniques for utilizing time and causality in determining the probability of a disease. We have also developed search techniques that allow NESTOR to explore efficiently a very large search space in order to find the most probable (multiple disease) hypothesis. This technique is general and can be applied to many nonmedical problems where the goal is to find the most probable hypothesis among many possibilities.

INTEGRATING MATHEMATICAL MODELS WITH AI METHODS

The work to date has concentrated on analyzing physiological function using knowledge of anatomy, physiology and causality. During the coming months, the causal analysis will be generalized. The issue is how to make valid qualitative conclusions based upon the quantitative relations of a physiological model. Rules will be developed and tested for inferring qualitative relations such as the following:

--inferring the direction of change in a dependent parameter following change in an independent parameter;

--inferring whether qualitatively significant change in a dependent parameter may occur as a consequence of a qualitatively significant change in an independent parameter; and

--inferring that a dependent parameter has a qualitatively significant value.

Questions of the following type will build on the other structures, incorporating the kinds of mathematical descriptions of human physiology necessary to predict the effects of quantitative deviations from normal:

"What is the effect of infusing hyperosmotic saline? How much should be used to return plasma volume to normal? How fast will normalization occur?"

PSYCHOLOGICAL STUDIES OF EXPLANATION

Within the next several months, Mr. Teach will have completed all data analysis for his research and will have completed most of his dissertation writing. Thereafter we will work on analyzing the results in light of our interest in the design of explanation capabilities for medical expert systems. We expect to have several papers based on this work completed or in preparation by this time next year.

EXPLANATIONS FOR A RULE-BASED EXPERT SYSTEM

As described in the GUIDON report elsewhere in this document, Dr. Rennels and others from the NEOMYCIN project will continue in their effort to implement a system to provide explanations for that program. This project will allow us to learn about optimal techniques for encoding and explaining the strategic tasking knowledge that is interfaced with that system's domain knowledge.

B. Requirements for Continued SUMEX Use

All the work we are doing is totally dependent on the SUMEX resource. As of yet, none of these young projects has been sufficiently mature to justify their transfer to one of the SUMEX personal workstations, so the KI-10's continue to be key elements in our research plan.

In addition, we have long appreciated the benefits of GUEST and network access to the programs we are developing. SUMEX greatly enhances our ability to obtain feedback from interested physicians and computer scientists around the country. As our programs mature in the months ahead, it will become increasingly important that we be able to make them available for demonstration and for access by distant collaborators via the SUMEX network.

C. Requirements for Additional Computing Resources

The mainframe machine should continue to provide a suitable environment for our work in the months ahead. We have no plans to transfer to other hardware soon. However, the response time on the main machine continues to be a major problem, even though the SUMEX workstations have provided additional cycles to permit off-loading of some users from the PDP-10.

D. Recommendations for Future Community and Resource Development

We concur with the suggestion, expressed by some of the other project reports, that consideration should be given to an upgrade of SUMEX to a more modern central mainframe machine.

II.A.1.5 MOLGEN Project

MOLGEN - A Computer Science Application to Molecular Biology

Prof. E. Feigenbaum and Dr. P. Friedland
Department of Computer Science
Stanford University

Assoc. Prof. D. Brutlag
Department of Biochemistry
Stanford University

Assoc. Prof. L. Kedes
Department of Medicine
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The MOLGEN project has focused on research into the applications of symbolic computation and inference to the field of molecular biology. This has taken the specific form of systems which provide assistance to the experimental scientist in various tasks, the most important of which have been the design of complex experiment plans and the analysis of nucleic acid sequences. We plan to expand and improve these systems and build new ones to meet the rapidly growing needs of the domain of recombinant DNA technology. We do this with the view of including the widest possible national user community through the facilities available on the SUMEX-AIM computer resource.

It is only within the last few years that the domain of molecular biology has needed automated methods for experimental assistance. The advent of rapid DNA cloning and sequencing methods has had an explosive effect on the amount of data that can be most readily represented and analyzed by computer. Moreover we have already reached a point where progress in the analysis of the information in DNA sequences is being limited by the combinatorics of the various types of analytical comparison methods available. The application of judicious rules for the detection of profitable directions of analysis and for pruning those which obviously lack merit will have an autocatalytic effect on this field in the immediate future.

The MOLGEN project has continuing computer science goals of exploring issues of knowledge representation, problem-solving, and planning within a real and complex domain. The project operates in a framework of collaboration between the Heuristic Programming Project (HPP) in the Computer Science Department and various domain experts in the departments of Biochemistry, Medicine, and Genetics. It draws from the experience of several other projects in the HPP which deal with applications of artificial intelligence to medicine, organic chemistry, and engineering.

We have begun a transition from being primarily a computer science research project to being an interdisciplinary project with a strong applications focus. The tools that we have already developed will be improved to the point where they make a significant contribution to both research and engineering in the domain of molecular biology.

B. Medical Relevance and Collaboration

The field of molecular biology is nearing the point where the results of current research will have immediate and important application to the pharmaceutical and chemical industries. Already, clinical testing has begun with synthetic interferon and human growth hormone produced by recombinant DNA technology. Governmental reports estimate that there are more than 200 new and established industrial firms already undertaking product development using these new genetic tools.

The programs being developed in the MOLGEN project have already proven useful and important to a considerable number of molecular biologists. Currently several dozen researchers in various laboratories at Stanford (Prof. Paul Berg's, Prof. Stanley Cohen's, Prof. Laurence Kedes', Prof. Douglas Brutlag's, Prof. Henry Kaplan's, and Prof. Douglas Wallace's) and over 400 others throughout the country are using MOLGEN programs over the SUMEX-AIM facility. We have exported some of our programs to users outside the range of our computer network (University of Geneva [Switzerland], Imperial Cancer Research Fund [England], and European Molecular Biology Institute [Heidelberg] are examples).

C. Highlights of Research Progress

Accomplishments:

The current year has seen intensive work on what might be termed a second-generation experiment design system. We have also begun to explore the capabilities of the SUMEX Dolphin computers for MOLGEN work. This section will summarize those projects as well as continuing work on other MOLGEN research.

1. Representation Research

The domain of molecular biology has proven a fruitful testbed in the development of a flexible software package, the Unit System, for symbolic representation of knowledge. The package is already in use by a variety of research projects both within the Heuristic Programming Project at Stanford and at other institutions. It provides for acquisition and storage of many different types of knowledge, ranging from simple declarative types like integers and strings to complex declarative types like nucleic acid restriction maps to procedural types like a rule language in a subset of English.

However, it has become clear that there are several bottlenecks in practical use of the Unit System on the main SUMEX computer system. First, there is the address space limitation. It is well known that Interlisp

programs on the DEC 10 and 20 series of computers have relatively little user memory available. Several of the MOLGEN knowledge bases have already exceeded that limitation. Second, it is clear that communication with the Unit System via 1200 baud display terminals does not allow the presentation of enough information at a fast enough speed. We are very optimistic that both of these problems will be solved by use of the Xerox Dolphin workstation computer. The Dolphin has a far larger address space than DEC 10 and 20 computers, as well as a large bit-map display with window and graphic capabilities.

Our initial experiments on the Dolphin have been encouraging. We were able to move the Unit System to the Dolphin with almost no difficulties. We have implemented a simple display interface that allows far more information to be directly accessible to the knowledge base builder. We are now in the state of allowing several MOLGEN biologists to begin using the Dolphins as practical knowledge entry tools.

2. Planning Research

The problem of designing laboratory experiments in molecular biology has been fundamental to MOLGEN research. Previous work consisted of two major subparts, each resulting in a doctoral thesis in computer science. The two systems, developed by Peter Friedland and Mark Stefik, produce reasonable experiment designs on test problems suggested by laboratory scientists.

The last year has seen the design and construction of a hybrid planning system combining the best of the previous two experiment planning systems. This work was done by Yumi Iwasaki, a master's candidate in computer science. The system, known as SPEX, uses Friedland's method of experiment design by the stepwise refinement of abstracted or "skeletal" plans, and Stefik's method for planning strategy control by "layers of planning spaces." It also keeps extensive records of all decisions made during the planning process; this will allow SPEX to be used as an experimental laboratory for work on explanation, verification, optimization, and debugging of plans.

In addition to the above general planning work, Rene' Bach has designed and implemented a small expert system for the problem of suggesting sequencing strategies for large nucleic acid molecules. The system, MAXAMIZE, was built entirely within the Unit System. It takes information about a given molecule (restriction map, already sequenced regions, etc.) along with information specific to the laboratory (how many bases may be comfortably sequenced by that lab, enzymes available, etc.) and provides an optimal or near-optimal strategy.

3. Knowledge Base Construction

Over seven man-years have now been spent in constructing knowledge bases for various fields of molecular biology. A new project was initiated this year with Professors Stanley Falkow and Esther Lederberg of the Stanford Medical Microbiology Department. Also, an entire series of

knowledge-base development projects has begun at the Imperial Cancer Research Fund in London, England, using the Unit System which was exported to a DEC 2060 at the ICRF.

The knowledge bases developed under the MOLGEN grant have begun to find their way into the daily laboratory practice of many of the scientists associated with the project. They have a provided a mechanism for managing the explosive growth of data and strategies in many areas of molecular biology without the necessity of building special purpose systems for each area. Also, the expert scientists themselves have been able to design and build their own systems, avoiding the time and reliability problem of a knowledge base passing through the filter of a computer scientist intermediary. The knowledge bases have served as "intelligent encyclopedias," as simulation systems, and as training vehicles.

It should also be noted that the Unit System allows for the easy transfer of knowledge from one knowledge base to another, and indeed the various expert molecular biologists have freely shared information as they work on related knowledge bases.

4. Other Applications of Symbolic Computation to Molecular Biology

MOLGEN work in the previous year has concentrated on basic research, there now being adequate commercial avenues for applied development of symbolic computation tools for molecular biology. However, we have continued to support large-scale guest access to the SEQ program for nucleotide sequence analysis and the Los Alamos Nucleic Acid Sequence Library. The GENET guest community has continued to grow. Because of this growth, a committee, headed by Professor Allan Maxam of Harvard University, has been formed to set policy for future guest access via SUMEX or other possible nationally networked systems.

We believe that GENET has been a clear demonstration that SUMEX is a model for building national communities of research scientists. We hope that the GENET committee will find adequate mechanisms for supporting this community without detracting from the primary research goals of SUMEX-AIM.

Research in Progress:

The next year will be spent in testing the SPEX experiment design system in the domain of cloning experiments, continuing the testing of the Dolphins for knowledge acquisition, and starting to explore the areas of plan debugging and explanation. We are also beginning work on automatic acquisition of experiment design strategies and discovery of gene regulation signals in the nucleic acid sequence library.

1. Cloning Advisory System

The SPEX experiment design system is now operational. In parallel with its development, we have begun to construct a large knowledge base specific to tools and techniques used during cloning experiments. The basic strategy or skeletal plan of all cloning experiments is quite

straight-forward: First, isolate the piece of DNA you wish to clone, second, select a vector to carry the clone, third, insert the DNA into the vector, fourth, select a host for expression of the hybrid molecule, fifth, insert the hybrid into the host, and sixth, select for the protein or nucleic acid product that was the eventual goal of the cloning experiment. Following this skeletal plan, the cloning knowledge base will contain information on DNA isolation methods, cloning vectors, insertion methods, hosts, host insertion methods, and selection methods. Professor Stanley Falkow and several others are assisting us in the building of this expert knowledge base.

We expect the cloning knowledge base to be essentially complete by the end of this year. We will also have accumulated substantial experience with running SPEX on that knowledge base. Sometime during this year, we will release the cloning advisory system to several of our molecular biologist colleagues for further testing and improvement. We also hope to gain more information about experiment design heuristics in general, using SPEX as a laboratory for this testing. The issues to be studied include optimal methods for plan expansion: depth-first breadth-first, and heuristic are among the possibilities, and how abstract or detailed to make the skeletal plans which drive the experiment design system.

2. Research on Plan Debugging and Explanation

SPEX keeps complete records of all decisions made during the course of designing an experiment. These include strategic decisions as to which general planning heuristics to employ and which domain-specific skeletal plans to use, as well as tactical decisions made in the course of choosing specific operators to instantiate a plan step. In addition, SPEX keeps a dynamic model of the world state as assumed after the execution of each plan step.

We will now make use of this information to add plan explanation and debugging facilities to the SPEX system. Users of SPEX will be able to easily determine which strategies and rules led to each planning decision. A limited querying interface will be added, probably more graphical through the Dolphin facilities than natural language.

Plan debugging can take two general forms. First, the debugging of plans generated originally by SPEX. In that case the debugging process will consist of attempting to ascertain the point at which the plan failed by comparing the actual state of the experiment at various points with SPEX's model of the world at those points. Presumably this will lead to identifying the instantiation step which failed; if no step failed, then the skeletal plan which generated the experiment design must have been faulty. In the former case, an alternate instantiation will be suggested (and possibly the knowledge base updated to reflect the error). In the latter case, an alternate skeletal plan will be used and the faulty one will be corrected (initially manually, but we hope by some automated means in the future).

The second general form of plan debugging occurs when the plan was generated outside of SPEX. In this case a plan verification step must first take place to ensure that the original plan design and model correspond to the knowledge base. If errors are found, then alternatives will be suggested, if not then debugging will proceed as above.

3. Further Developments of the Unit System

Considerable effort will be spent in taking full advantage of the many capabilities for graphical interaction of the Dolphin. We expect improvements to be opportunistic as our molecular biologist collaborators begin to do their routine knowledge base development on the Dolphin.

4. Work in Plan and Model Discovery

Initial exploration has begun in two new projects in what might be termed automated knowledge acquisition. The first has to do with acquiring skeletal plans for the SPEX knowledge base. Currently skeletal plans are described directly by molecular biologists. We would like to explore means of automatically extracting these abstracted plans from successful experiment designs, either from the literature or in an interactive mode from scientists. One of the major questions to be resolved concerns determining how specific or abstract to make the skeletal plans; specificity leads to efficiency, abstractness leads to generality--a balance is needed.

The second project is in the automated discovery of rules for the regulation of genes from the data stored in the nucleotide sequence library. It is well known that close to 90% of all DNA in higher organisms does not code directly for proteins, but has to do with the regulation of those genes that do code for proteins. There is currently no widely believed mechanism which explains even an appreciable fraction of regulatory processes. The information used for gene regulation is stored somewhere within the nucleotide sequence itself. We have access to a large collection of such sequence data (over 600000 bases) and to the best currently available sequence analysis system (SEQ). We plan to add semantic pattern finding procedures to SEQ, probably through a Unit System interface, that will allow proposed theories to be tested on the data base.

D. Publications

Bach R., Friedland P., Brutlag D., Kedes L., MAXAMIZE, A DNA Sequencing Strategy Advisor, Nucleic Acids Research, Volume 10, Number 1, 295-304 (January 1982) Brutlag D., Clayton J., Friedland P., Kedes L., SEQ: a Nucleotide Sequence Analysis and Recombination System, Nucleic Acids Research, Volume 10, Number 1, 279-294 (January 1982)

Clayton J., Kedes L., GEL, a DNA Sequencing Project Management System, Nucleic Acids Research, Volume 10, Number 1, 305-321 (January 1982)

Feitelson J., Stefik M.J., A Case Study of the Reasoning in a Genetics Experiment, Heuristic Programming Project Report HPP-77-18 (Working Paper) (May 1977)

Friedland P., Knowledge-Based Experiment Design in Molecular Genetics, Proceedings Sixth International Joint Conference on Artificial Intelligence, 285-287 (August 1979)

Friedland P., Knowledge-Based Experiment Design in Molecular Genetics, Ph.D. Thesis, Stanford CS Report CS79-760 (December 1979)

Friedland P., Kedes L., Brutlag D., MOLGEN--Applications of Symbolic Computation and Artificial Intelligence to Molecular Biology, Proceeding of the Battelle Conference on Genetic Engineering (April 1981)

Friedland P., Acquisition of Procedural Knowledge from Domain Experts, Proceedings Seventh International Joint Conference on Artificial Intelligence, 856-861 (August 1981)

Friedland P., Kedes L., Brutlag D., Iwasaki Y., Bach R., GENESIS, a Knowledge-Based Genetic Engineering Simulation System for Representation of Genetic Data and Experiment Planning, Nucleic Acids Research, Volume 10, Number 1, 323-340 (January 1982)

Iwasaki Y., and Friedland P., SPEX: A Second-Generation Experiment Design System, to appear in Proceedings of the Second National Conference on Artificial Intelligence (August 1982).

Martin N., Friedland P., King J., Stefik M.J., Knowledge Base Management for Experiment Planning in Molecular Genetics, Fifth International Joint Conference on Artificial Intelligence. 882-887 (August 1977)

Stefik M., Friedland P., Machine Inference for Molecular Genetics: Methods and Applications, Proceedings of the National Computer Conference, (June 1978)

Stefik M.J., Martin N., A Review of Knowledge Based Problem Solving As a Basis for a Genetics Experiment Designing System, Stanford Computer Science Department Report STAN-CS-77-596. (March 1977)

Stefik M., Inferring DNA Structures From Segmentation Data: A Case Study, Artificial Intelligence 11, 85-114 (December 1977)

Stefik, M., An Examination of a Frame-Structured Representation System, Proceedings Sixth International Joint Conference on Artificial Intelligence, 844-852 (August 1979)

Stefik, M., Planning with Constraints, Ph.D. Thesis, Stanford CS Report CS80-784 (March 1980)

E. Funding Support

The MOLGEN grant is titled: MOLGEN: A Computer Science Application to Molecular Biology. It is NSF Grant ECS-8016247. Current Principal Investigators are Edward A. Feigenbaum and Bruce G. Buchanan, Professors of

Computer Science, Laurence H. Kedes, Investigator, Howard Hughes Medical Institute and Associate Professor of Medicine, and Douglas L. Brutlag, Associate Professor of Biochemistry. MOLGEN is currently funded from 10/81 to 9/82 at $159,785 including indirect costs as the second year of a three year renewal.

II. INTERACTIONS WITH THE SUMEX-AIM RESOURCE

SUMEX-AIM continues to provide the bulk of our computing resources. The facility has not only provided excellent support for our programming efforts but has served as a major communication link among members of the project. Systems available on SUMEX-AIM such as INTERLISP, TV-EDIT, and BULLETIN BOARD have made possible the project's programming, documentation and communication efforts. The interactive environment of the facility is especially important in this type of project development.

Unfortunately, the computing environment at SUMEX has suffered in the recent past from heavy demands on cycle time creating serious real-time delays for programmers and knowledge-base building especially. The Units Editor is especially sensitive because of its relatively large demands on cpu-resources. Accordingly, a significant fraction of the MOLGEN group activity has been transferred to the SCORE computer in the Department of Computer Science at Stanford. When SUMEX hardware is updated, we anticipate that its response time will improve and the MOLGEN computing will return full time to SUMEX. It is clear, however, that the MOLGEN project continues to thrive and prosper because of the computing environment only available at SUMEX: the interactive environment including instantaneous communications among collaborators who are physically distant (even on the Stanford campus), and especially the unique telecommunications facilities that have allowed the development of the GENET community with its access to MOLGEN applications tools are two clear examples.

We have taken advantage of the collective expertise on medically-oriented knowledge-based systems of the other SUMEX-AIM projects. In addition to especially close ties with other projects at Stanford, we have greatly benefitted by interaction with other projects at yearly meetings and through exchange of working papers and ideas over the system.

The ability for instant communication with a large number of experts in this field has been a determining factor in the success of the MOLGEN project. It has made possible the near instantaneous dissemination of MOLGEN systems to a host of experimental users in laboratories across the country. The wide-ranging input from these users has greatly improved the general utility of our project.

We find it very difficult to find fault with any aspect of the SUMEX resource management. It has made it easy for us to expand our user group, to give demonstrations (through the 20/20 adjunct system), and to disseminate software to non-SUMEX users overseas.

III. RESEARCH PLANS

A. Justification and Requirements for Continued SUMEX Use

The MOLGEN project depends heavily on the SUMEX facility. We have already developed several useful tools on the facility and are continuing research toward applying the methods of artificial intelligence to the field of molecular biology. The community of potential users is growing nearly exponentially as researchers from most of the bio-medical fields become interested in the technology of recombinant DNA. We believe the MOLGEN work is already important to this growing community and will continue to be important. The evidence for this is an already large list of pilot exo-MOLGEN users on SUMEX.

SUMEX is currently having difficulty meeting the research needs of the MOLGEN project adequately. We expect to need more file space as our knowledge bases grow; perhaps an additional 5000 disk blocks in the next few years for that work. Our real difficulties will come in the applications testing of MOLGEN tools. We support with great enthusiasm the acquisition of satellite computers for technology transfer and hope that the SUMEX staff continue to develop and support these systems. One of the oft-mentioned problems of artificial intelligence research is exactly the problem of taking prototypical systems and applying them to real problems. SUMEX gives the MOLGEN project a chance to conquer that problem and potentially supply scientific computing resources to a national audience of bio-medical research scientists.

II.A.1.6 MYCIN Projects Group

MYCIN Projects

Edward H. Shortliffe, M.D., Ph.D.
Departments of Medicine and Computer Science
Stanford University

Bruce G. Buchanan, Ph.D.
Computer Science Department
Stanford University

I. SUMMARY OF RESEARCH PROGRAM

A. Project Rationale

The MYCIN Projects are a set of related research programs, each devoted to the development of knowledge-based expert systems for application to medicine and the allied sciences. The name was derived from our first system, the MYCIN program. That research has now given way to three sub-projects (EMYCIN, GUIDON, and ONCOCIN), each of which is discussed in the sections that appear below. The key issue for all sub-projects has been to develop programs that can provide advice similar in quality to that given by human experts, and to develop systems that are easy to use and acceptable to physicians and medical students.

The success of the original MYCIN infectious disease consultation program has led us to try to generalize and expand the methods employed in that program to a number of ends:

(1) to develop consultation systems for other domains (our generalized system-building tool is known as "Essential MYCIN", or EMYCIN, and has been applied in several new areas; ONCOCIN is our a more recent consultation system and was inspired by EMYCIN although it is an entirely new program);

(2) to explore other uses of the MYCIN knowledge base (our tutoring system, GUIDON, uses the infectious disease knowledge in MYCIN to teach medical students about diagnosis and management of infections);

(3) to continue to improve the interactive process, both for the developer of a knowledge-based system, and for the user of such a system (both EMYCIN and ONCOCIN have stressed simplified techniques for interacting with a knowledge base and entering data); and

(4) to experiment with alternate techniques for knowledge representation, recognizing that the pure production rule method used in MYCIN was inadequate at times and frequently led to confusion regarding the separation of strategy or control processes from domain knowledge (ONCOCIN uses production rules as only one of several knowledge

representation techniques, and the work on GUIDON has led to a more robust revised version of MYCIN known as NEOMYCIN).

B. Medical Relevance and Collaboration

By utilizing our EMYCIN system to collaborate on building the PUFF program, we learned that it is possible in a short period of time to develop a clinically useful consultation system using the domain-independent parts of MYCIN. EMYCIN has since been applied in a number of additional medical domains. With each successive application we learn more about the representation of medical knowledge and the scope and limitations of the production rule formalism used in EMYCIN. For example, it has become clear that "shallow" rules relating signs and symptoms to diagnoses through a few intermediate concepts can be sufficient for high performance in medical diagnosis. On the other hand, such shallow rules are not always sufficient for teaching medical students because they lack the deeper causal links needed for justifying and remembering the more shallow associations (see GUIDON discussion below).

Although EMYCIN was not used to build our new ONCOCIN program, the lessons learned in building prior production rule systems have allowed us to create a large oncology protocol management system much more rapidly than was the case when we started to build MYCIN. We introduced ONCOCIN for use by Stanford oncologists in May 1981. This would not have been possible without the active collaboration of Stanford oncologists who helped with the construction of the knowledge base and also kept project computer scientists aware of the psychological and logistical issues related to the operation of a busy outpatient clinic.

In addition, there is a growing realization that medical knowledge, originally codified for the purpose of computer-based consultations, may be utilized in additional ways that are medically relevant. Using the knowledge to teach medical students is perhaps foremost among these, and GUIDON continues to focus on methods for augmenting clinical knowledge in order to facilitate its use in a tutorial setting. A particularly important aspect of this work is the insight that has been gained regarding the need to structure knowledge differently, and in more detail, when it is being used for different purposes (e.g., teaching as opposed to clinical decision making). This aspect of the GUIDON research has led to the development of a modified version of MYCIN, NEOMYCIN, which is an evolving computational model of medical diagnostic reasoning that we hope will enable us to better understand and teach diagnosis to students.

C. Highlights of Research Progress

1. Accomplishments This Past Year

EMYCIN

In the last year, work on EMYCIN was brought to a conclusion. The program was "frozen" and documentation completed. The program remains a

vehicle for computer science research, but there is no funding for continued development of the tool. Our work centered around the following activities:

(1) One of the most significant experiments with the program is the development of ROGET, a knowledge acquisition program that aids a system builder with the definition and creation of a new knowledge base for an EMYCIN system. ROGET itself is written in EMYCIN. Design was completed in this year and programming started, as described in more detail below.

(2) We continued to help users in the SUMEX community with new applications of EMYCIN and with questions about on-going applications. These are described below.

(3) Buchanan and Shortliffe have begun writing up the results of many years' work on MYCIN and EMYCIN. Because we view AI as an experimental science, we believe there is much to be learned in a detailed analysis of the experiments performed with rule-based consultation systems by the MYCIN group. This work is also briefly described below.

ROGET: A KNOWLEDGE-BASED SYSTEM FOR THE DESIGN OF EXPERT SYSTEMS

ROGET is a program under construction which is designed to assist an expert with the initial formulation of an EMYCIN consultant. The ROGET system is intended (1) to help the expert properly delimit the scope and function of a consultant, (2) to identify the major types of inferences it will make and their requisite facts, and (3) to educate the expert about the use of the particular representation system being used to construct the consultant.

Current expert-system building tools provide little or no guidance to an expert during the initial formulation of a consultant system. A knowledge engineer provides the expert with advice about what tasks the consultant might have to perform and, more importantly, how to employ a particular representation system to achieve this performance. Previous work by Davis, Reboh, and van Melle have emphasized tools and techniques that are appropriate only considerably after the basic "inference structure" of the consultant has already been designed and placed on-line. These tools are meant to assist the refinement of a knowledge base and they make the implicit assumption that the expert is already fairly familiar with (i.e., has at least reading knowledge of) the representation system.

In contrast, the primary function of ROGET will be to guide and instruct the expert as he selects the relevant concepts and inferences required to solve a problem, indicate how they have been typically represented within a specific representational framework, and provide hands-on instruction regarding the tool's actual use. ROGET will provide an interactive environment in which the expert can construct a prototypical inference structure for a consultant and some of its "connective tissue". In this manner, the expert can become familiar enough with the

representation system to use the existing knowledge acquisition tools to further refine the knowledge base.

This research is predicated on the observation that several computer-based consultants, particularly of the diagnostic variety, employ similar types of inference, provide similar types of advice, and require similar types of data. Furthermore, the actual representation of the domain-specific counterparts of this knowledge within a particular representation system has also taken rather stereotypic forms. This research will attempt to capitalize on these "standard" forms by representing this type of knowledge and developing a system that will use it to guide the acquisition of new expertise in different domains.

Helping the expert identify an appropriate task for the consultation system is a major function of any knowledge engineer. ROGET will assist the expert with the proper scoping of the consultant by identifying the applicable inferences for this problem and by suggesting the most appropriate inferences for the current consultant to make. These suggestions can be made on the basis of information about the experts background, his experience with system construction, any constraints on his time and other resources, and, most importantly, a description of the "consultation setting", where and how the consultant will interact with prospective clients.

The education of the expert about the intricacies of a particular representation system is currently accomplished primarily through the use of examples taken from other expert systems constructed with this tool. ROGET will mimic this behavior as well. In addition to providing domain-specific examples of each of the inference, advice, and data types mentioned above, ROGET will also give examples of how the inference engine interprets example knowledge bases. Again, knowing about the consultation setting will help ROGET focus attention on specific portions of complete knowledge bases that demonstrate how certain types of knowledge are needed to make certain features of the inference engine run.

ROGET will be developed to assist an expert design an EMYCIN consultant. To represent the consultation setting a simple, script-like representation of the major events, people, and information flow for various consultation scenarios will be developed. In the case of diagnostic consultants this representation must be adequate to capture most "clinical settings" at some fairly high level of abstraction. The expert will then identify specific types of people, information, and events that are applicable to his domain and use this perspective to identify points in the information flow where a consultant would be most appropriate and effective.

This same script-like representation will also be used to capture the stereotypic inference steps that EMYCIN diagnostic consultants employ, to identify what types of data are expected at certain points in the consultation, and to understand how the advice the consultant gives will be used by the client. The current inference engine of EMYCIN will be modified to operate on the basis of this new representation of control knowledge, in

addition to the normal backward-chaining of rules. Furthermore, the explanation machinery will be updated to provide alternative explanations of the consultant's line of reasoning about a case based on this representation of control.

Representation of this scriptal knowledge, the hierarchies of inferences, data, and advice for each task, and the examples themselves will employ a simple UNITS-style representation system embedded within the current EMYCIN system. A trace of the design process for individual consultants will be kept and used as the basis for "developmental explanations" that could be used to indicate why the consultant is scoped to perform certain tasks to the exclusion of others and to identify parts of the knowledge base that are under-developed.

The initial ROGET system will be considered "done" when it can successfully scope and design the basic inference structures for the following EMYCIN consultants: MYCIN, GRAVIDA, PUFF, and SACON. At least one general consultant task must be represented; currently this task is diagnosis and simple therapy selection. Another potential task would be the design of objects or plans by configuring schematic designs (ROGET itself, R1, and Friedland's MOLGEN are examples of this type of "configuration" consultant).

In summary, then, this project seeks to make explicit some important aspects of knowledge-engineering expertise. It will identify a taxonomy of the kinds of expert systems in terms of common tasks and how their place in a larger information context effects their final design. This knowledge will be used to help scope the initial design of a consultant and to educate the expert about a representation system. It is concerned with identifying expertise required for initial expert system design and this, in turn, will require an analysis of the nature of expertise itself and the current limitations on its representation, especially within the EMYCIN system.

EMYCIN Applications

We spoke with many physicians around the country about potential applications of EMYCIN in their own work. We encouraged several to submit applications to SUMEX for pilot status. We collaborated with one physician at length on a prototype system named COAGS, described below.

COAGS

During January and February of 1982 James Bennett and Robert T. Wall, M.D., a hematologist at the Stanford Medical School, collaborated to construct a prototype hematology consultant using the EMYCIN system. The consultant, named COAGS, deals with patients about to undergo surgery when routine coagulation screening tests indicate a potential bleeding problem. The surgeon will typically refer the patient to a hematologist for a more complete set of tests to confirm the disorder and to receive a recommendation about any treatment that might allow the operation to proceed.

The consultant requests that various collections of tests be performed in an attempt to confirm that a problem does exist. Most of the coagulation tests are not particularly accurate and repetition of critical tests with appropriate controls is an essential part of hematologists' attempts at confirmation. The tests can be expensive to perform and the consultant has been constructed to minimize the cost of diagnosing the most specific bleeding disorder possible.

The EMYCIN system did not require any improvements or modifications to represent Dr. Wall's clinical expertise. The consultant was constructed in a two month time period on weekly meetings between expert and knowledge engineer. We believe that this represented approximately 6 person-days.

ANALYSIS OF EXPERIMENTS WITH MYCIN & EMYCIN

Artificial Intelligence, or AI, is largely an experimental science in that at least as much progress has been made by building and analyzing programs as by examining theoretical questions. MYCIN is one of several well-known programs that embody some intelligence and provide data on the extent to which intelligent behavior can be programmed. As with other AI programs, its development was slow and not always in the forward direction. But we feel we learned some lessons in the course of over eight years of work on MYCIN, and related programs. The purpose of this analysis, then, is to share the results of many experiments performed in that time, and to paint a coherent picture of the work.

The MYCIN project is one of the clearest representatives of the experimental side of AI. It was begun in the spring of 1972 with a set of discussions among medical school and computer science researchers interested in applying more intelligence to computer programs that interpret medical data. Shortliffe's PhD dissertation in 1974 discusses the problem context and the MYCIN program that implemented a solution.

In itself, the 1974 version of MYCIN represents an experiment. We were testing the hypothesis, tentatively advanced in the DENDRAL work, that a production rule formalism was sufficient for the high performance, flexibility, and understandability that we demanded in an expert system. The positive answer to this question is one of the best-known lessons in the history of AI.

In addition to, or rather because of, the original MYCIN program and the knowledge of blood infections and meningitis that was accumulated for that work, many derivative projects explored variations on the original design. In a book-length analysis currently underway, we are examining many of the experiments that evolved in the period from 1972 to 1980 based on the 1972-74 design effort. We have chosen those pieces of work which, at least in retrospect, can be seen as posing clear questions and producing clear results that were documented in the AI or medical literature, or occasionally in technical reports.

Some of those other projects are discussed in this annual report, for example GUIDON and NEOMYCIN. Others have come to a logical conclusion:

TEIRESIAS, VM, CENTAUR, and now EMYCIN. In those major pieces of work, and several others as well, the questions posed and the answers we derived from program construction and testing are important scientific results. Documenting these results will, we hope, benefit the entire AI community.

Because almost all of the computing work was done on the SUMEX-AIM computer resource at Stanford, we recognize that this set of experiments would not have been possible without funding from the Division of Research Resources of the National Institutes of Health.

GUIDON

The purpose of the GUIDON project is to develop a tutorial program that uses a knowledge-based system for teaching material. Research over the past five years has shown that a knowledge representation that is adequate for developing a consultation system with good performance may be inadequate for teaching. Our study of MYCIN's meningitis rule base (1979-1980) led us to develop a new representation in which kinds of knowledge important to teaching are represented separately and explicitly. Using this new system, which we call NEOMYCIN, we proceeded in 1981 to develop a student modelling program that will be used in a tutorial program for teaching diagnostic strategies to students.

CONTINUING DEVELOPMENT OF NEOMYCIN

NEOMYCIN is distinguished from other AI consultation programs by its uses of an explicit set of domain-independent meta-rules for controlling all reasoning. These rules constitute the diagnostic procedure that we want to teach to students: the stages of diagnosis, how to focus on new hypotheses, and how to evaluate hypotheses. It has been a major undertaking, separate from the problem of representing disease knowledge, to design and test this diagnostic procedure. For example, this year we extracted the forward-directed reasoning that had previously been embedded in NEOMYCIN's interpreter and represented this as a "task" with 4 meta-rules (one meta-rule each for: follow-up questions, antecedent rules, trigger (hypothesis-suggesting) rules, and forward application of rules that confirm hypotheses already under consideration). By the end of the year, attention turned to extending NEOMYCIN's disease knowledge to include more diseases that present similarly to, or are confused with, meningitis.

THE IMAGE STUDENT MODELLING PROGRAM

Teaching diagnosis involves recognizing the intent behind a student's behavior, so that missing knowledge can be distinguished from inappropriate strategies. The teacher interprets behavior, critiques it, and provides advice about other approaches. To do this successfully and efficiently in a complex domain, the teacher benefits from multiple, complementary modeling strategies. IMAGE is a student modelling program that uses NEOMYCIN's meta-rules and disease knowledge to understands student diagnostic plans by a dual search strategy. IMAGE first produces multiple predictions of student behavior by a model-driven simulation of NEOMYCIN. Focused, data-driven searches then explain incongruities. By supplementing

each other, these methods lead to an efficient and robust plan understander.

A model of student strategies in medical diagnosis must disambiguate the possible purposes and knowledge underlying the student's actions. The approaches followed by other plan recognizers and student modelers are not sufficient here because:

(1) the complex domain makes thorough searches impractical, whether top-down or bottom-up;

(2) we are not modeling only facts and rules used in isolation, but also the procedures for applying them;

(3) every one of the student's actions must be monitored in case the teaching module decides to interrupt;

(4) his behavior must be evaluated and not just explained; and

(5) we might not have any explicit goal statements from the student, so we expect to rely only on his queries for problem data as evidence for his thinking.

The IMAGE program is now operational and will soon be tested in experiments with medical students.

SHORT-TERM PLANS FOR NEOMYCIN AND GUIDON

The next stages in the development of the GUIDON2 tutorial program (using IMAGE and NEOMYCIN) are:

(1) additions to NEOMYCIN's disease knowledge so we can fairly evaluate the program's focussing strategies;

(2) experimentation to determine the capabilities of IMAGE for explaining student (non-expert) behavior;

(3) development of an explanation system for NEOMYCIN, to serve as a testbed for explanations to be generated by the tutorial program (including use of a student/user model and condensation methods for presenting explanations at the appropriate level of detail) (Drs. Clancey and Shortliffe, and their graduate students, are collaborating on this project);

(4) design of various teaching scenarios for teaching NEOMYCIN's meta-rules, and representation of this teaching knowledge in a prototype version of GUIDON2.

ONCOCIN

The oncology chemotherapy consultation system, named ONCOCIN, has achieved many of its goals since work on the project began in July 1979. We are developing an interactive system to be used by oncology faculty and fellows in the Debbie Probst Oncology Day Care Center at Stanford University Medical Center.

The ONCOCIN research goals are directed both towards the basic science of artificial intelligence and towards the development of clinically useful oncology consultation tools.

We have undertaken AI research with the following aims:

(1) to implement and evaluate recently developed techniques designed to make computer technology more natural and acceptable to physicians [formal studies are underway to assess the success of our prototype system in achieving these goals];

(2) to extend the methods of rule-based consultation systems to interact with a large database of clinical information; and

(3) to continue basic research into the following problem areas: mechanisms for handling time relationships, techniques for quantifying uncertainty and interfacing such measures with a production rule methodology, approaches to acquiring knowledge interactively from clinical experts, assessment of knowledge base completeness and consistency.

The clinic system itself is intended to have a number of performance capabilities when it is completed. We outline those goals here, citing those which have been completely or partially accomplished and indicating those that have yet to be achieved:

(1) to assist with identification of current protocols that may apply to a given patient [not yet undertaken; will not be relevant until more protocols than lymphomas have been encoded]

(2) to assist with determining a patient's eligibility for a given protocol [not yet undertaken because lymphoma patients have been fully screened before they come to chemotherapy clinic];

(3) to provide detailed information on protocols in response to questions from clinic personnel [a query system is just now being started; it is being designed to assist both physicians and those building a protocol knowledge base];

(4) to assist with chemotherapy dose selection and attenuation for a given patient [fully implemented and being evaluated for patients under treatment for lymphoma];

(5) to provide reminders, at appropriate intervals, of follow-up tests and films required by the protocol in which a given patient is enrolled [fully implemented and being evaluated for patients under treatment for lymphoma];

(6) to reason about managing current patients in light of stored data from previous visits of (a) the individual patients [partially achieved, but much work remains], or (b) the aggregate of all "similar" patients [not yet attempted].

In the early years of the project, we developed a prototype of the ONCOCIN consultation system, drawing from programs and capabilities developed for the EMYCIN system-building project. We also undertook a detailed analysis of the day-to-day activities of the Stanford Oncology Clinic in order to determine how to introduce ONCOCIN with minimal disruption of an operation which is already running smoothly. Subsequently we completed the development of a special interface program that responds to commands from a customized keypad. Software protocols were developed for achieving communication between the interface program and the reasoning program, and we coordinated the printing routines needed to produce hardcopy flowsheets, patient summaries, and encounter sheets. Finally, lines were installed in the Stanford Oncology Day Care Center, and, beginning in May 1981, eight fellows in oncology began using the system three mornings per week for management of their patients enrolled in lymphoma chemotherapy protocols.

The introduction of the system last year has largely determined the nature of our subsequent research effort. That experience has proved to be a constant source of new ideas and challenges for system improvement. Much of our work in the last year has been in response to the lessons learned by observing the strengths and weaknesses of the computing system now that it is in routine use by physicians. Formal evaluation studies have continued in parallel with the ongoing development effort, and the data collection for them is within three months of completion; we expect to begin data analysis before the end of June.

IMPLEMENTATION OF ONCOCIN IN THE STANFORD ONCOLOGY CLINIC

ONCOCIN was introduced smoothly into the Oncology Clinic in May and June of 1981. Each oncology fellow was given approximately 45 minutes of training in the system's organization and in use of the terminal interface. They found it to be extremely easy to learn, and additional help was seldom needed after the training session. During the first several months our data manager was routinely present in the clinic to assist with questions, but it became clear that this was not necessary and since January of this year we have simply asked the physicians to call our project office if a problem arises.

Our Clinical Specialist Dr. Carlson, himself a fellow in oncology, has regularly reported to us the informal reactions of the physicians to the system. We will soon have formal data regarding their attitudes (see Study 1 below), but the informal feedback and suggestions were important

because they led to the design and implementation of several additional features that have required a good deal of programmer time during the past year. For example, the communications between the Reasoner and Interviewer were streamlined for efficiency, and major changes were made to a special category of data entries known as "special questions". These frequently annoyed the physicians because they tended to interrupt the smooth flow of the dialogue with the machine. The revised version of the interface inserts these questions more benignly into the flowsheet data form.

Two major problems were identified that relate to knowledge representation in ONCOCIN and control of the reasoning process. One of these is the inability of the program to manage complex inferences that are dependent upon analyzing a patient's response to therapy over all previous visits to the clinic. This kind of analysis is typically ignored in the formal protocols but is routinely used by experienced oncologists and is, of course, a desirable capability for a consultation program in this area. The development of techniques for managing temporal reasoning is a major interest of our group and is of particular importance in dealing with expert decision making in a domain such as cancer chemotherapy. Our oncology collaborators are finding this aspect of their knowledge particularly difficult to distill, and we expect that our attempts to develop computer-based approaches to the problem will provide particularly stimulating challenges in the months ahead.

The second problem was unanticipated until the system was implemented but is having a major impact on our future plans. ONCOCIN provides a therapy recommendation for the physician but allows those recommendations to be overridden if the physician prefers an alternate treatment. If the fellows change the therapy, they are asked to indicate the reason for preferring their therapy to what ONCOCIN suggested. When there is a major disagreement between the fellow and the computer, the physicians seldom object to indicating the reasons for the change, and their entries have been useful in further debugging the program's knowledge of lymphoma treatment. However, the changes they make are frequently minor ones that are more stylistic than substantive (e.g., some physicians prefer always to give prednisone in multiples of 20 mgs when they are dealing with large doses; others prefer to be more precise in their dosing and will ask patients to take combinations of 5 and 20 mg tablets). In these cases, the physicians tend to be annoyed by requests for reasons that they have disagreed with ONCOCIN; they do not see the slight adjustments as being real disagreements. We have attempted to respond to this problem by developing criteria for determining when a physician's preferred treatment is a clinically significant difference from what ONCOCIN has recommended. The knowledge for determining what changes are "important" is different from the kind of knowledge in the protocols themselves. It is more "judgmental" and imprecise, but seems to be an important addition to the program's knowledge base. Furthermore, this issue, and the observation that the physicians frequently believe they know what to do and would rather not be bothered by the program's suggestions, has led us to propose a modified version of ONCOCIN that will verify the physician's plan rather than offer a therapy recommendation for every case. We have referred to this new work as "hypothesis assessment" (see below).

EVALUATIONS OF THE ONCOCIN SYSTEM

All three ONCOCIN evaluations are approaching completion, and we expect to have data analyzed and formal reports written by this time next year.

Study I is an evaluation of the program's impact on the attitude of the oncology fellows towards computers in general and ONCOCIN in particular. All physicians were administered questionnaires and structured interviews in the Spring of 1981 before ONCOCIN was introduced. The same questionnaires have recently been distributed to them once again, and follow-up interviews are just now underway. Pre- and post-ONCOCIN results will be analyzed as described in last year's annual report.

Study II is an evaluation of the program's impact on the completeness and accuracy of flowsheet data recorded with and without ONCOCIN. Research programmers have written routines to formally analyze on-line flowsheets for completeness and accuracy. Pre-ONCOCIN flowsheets have been entered into the system exactly as they were originally recorded by the physician. The same analytic routines will analyze these pre-ONCOCIN flowsheets and we will compare the pre- and post-ONCOCIN results. We expect to begin analysis of the post-ONCOCIN data within the next two months when a sufficient number of patient visits have been recorded to permit a statistically significant comparison.

Finally, Study III is examining the comparison between ONCOCIN's therapeutic advice and the treatment decisions made by oncology fellows in the same setting. As described in last year's report, expert evaluators are rating treatment plans without knowing whether the recommendation is that of ONCOCIN or one of the clinic physicians. Over 200 flowsheets are currently being evaluated by Stanford lymphoma experts, and we expect to have their ratings returned for analysis by the end of May.

THE INTERVIEWER/REASONER MODEL

We have continued to refine ONCOCIN's communication system for permitting the physician to interact with a rapid interface program while a slower reasoning program considers treatment options in a second process. After the system was implemented, several unexpected problems resulting from the use of asynchronous processes were discovered and required resolution. For example, the reasoning program has no straightforward way of knowing whether a piece of information that it needs has been skipped on the flowsheet by the physician or whether the doctor has simply not reached that part of the flowsheet yet. Similarly, the Reasoner needs techniques for dealing with changes to flowsheet entries that are made after the initial entry has been transmitted by the Interviewer and has been used by the Reasoner in its decision making process.

HYPOTHESIS ASSESSMENT

As mentioned above, we have decided to consider modifications to ONCOCIN that would permit it to function as an "observer" of the physician's own decisions rather than as a primary source of advice. By permitting the physician to enter his or her own therapy plan on the flowsheet, we can acknowledge the oncologist's ability to reach appropriate therapeutic decisions for most patients. ONCOCIN will simply compare the physician's plan with what it believes is the proper therapy. If the system agrees with the physician, or determines that small differences are clinically insignificant, no advice from the computer will be necessary. If significant disagreements occur, on the other hand, ONCOCIN will need to respond with warnings and explanations for why it feels that an alternate therapy plan may be preferable. Our early experience with ONCOCIN since its clinic implementation suggests that this mode of interaction may be preferred by the clinic physicians. It will require minimal changes to ONCOCIN's decision making approach, but the determination of what differences are clinically significant, and the optimal method for explaining their importance to the physician, will be exciting challenges and important theoretical problems. This work is just getting underway and will receive much attention during the coming year.

QUERY SYSTEM

We have also recently started work on the development of a query system that will permit easy access to the large ONCOCIN knowledge base. Now that we have encoded several hundred rules, it is becoming unwieldy for system builders to work from large hard-copy listings of the knowledge base, and physicians will also require direct access to the program's knowledge. The query system is being designed to permit this kind of access. Rather than dealing with natural language understanding by computer, we are designing ways that menu selection and the high-speed interface can be used to permit access to the information that is needed by a physician or system builder.

VERIFYING THE COMPLETENESS AND CONSISTENCY OF THE KNOWLEDGE BASE

An important question for AI researchers involved with the development of expert systems is how to ascertain that a knowledge base for a consultation program is complete and consistent. Dr. Motoi Suwa, a visitor to Stanford from Japan two years ago, became fascinated with this question and collaborated with us on a formal analysis of the developing ONCOCIN knowledge base. However, the programs that he wrote were never formally linked to our system for writing rules and modifying other parts of the knowledge base. As a result, we have recently spent time modifying his code so that it will operate as an integral part of ONCOCIN. The development of an improved interface to permit this program to be used by our expert collaborators is also planned for the coming year.

Modeling: Prediction and Description. Submitted to AAAI-82.

D. Publications Since January 1981

1. Shortliffe, E.H. Medical consultation systems: designing for doctors. To appear in Communicating With Computers (M.J. Coombs, ed.), London: Academic Press, 1982.

2. Shortliffe, E. H. Medical Cybernetics: The Challenges of Clinical Computing. To appear in Technology and Growth, S. Basheer Ahmed, editor; Kennikat Press Corp., Port Washington, New York, 1982.

3. van Melle, W., Shortliffe, E. H., Buchanan, B. G. EMYCIN: A domain-independent system that aids in constructing knowledge-based consultation programs. In the Pergamon-Infotech State of the Art Report on Machine Intelligence, August 1981.

4. Shortliffe, E.H., Scott, A.C., Bischoff, M.B., Campbell, A.B., van Melle, W., and Jacobs, C.D. ONCOCIN: An expert system for oncology protocol management. Proceedings of 7th IJCAI, pp. 876-881, Vancouver, B.C., August 1981.

5. Clancey, W.J. and Letsinger, R. NEOMYCIN: reconfiguring a rule-based expert system for application to teaching. Memo HPP-81-2, Heuristic Programming Project, Stanford University, February 1981.

6. Suwa, M., Scott, A.C., and Shortliffe, E.H. An approach to verifying completeness and consistency in a rule-based expert system. Memo HPP-81-5, Stanford Heuristic Programming Project, April 1981; revised April 1982.

7. Teach, R.L. and Shortliffe, E.H. An analysis of physician attitudes regarding computer-based clinical consultation systems. Comput. Biomed. Res. 14:542-558 (1981).

8. Shortliffe, E.H. Expert systems research: adapting technical knowledge for computer-based consultation systems. Proceedings of the Sixth Symposium on Legal Data Processing, Sponsored by the Council of Europe, Thessaloniki, Greece, July 1981.

9. Shortliffe, E.H. Evaluating Expert Systems. Memo HPP-81-9, Stanford Heuristic Programming Project, July 1981. Contributed to Chapter 6 of Building Expert Systems, edited by R. Hayes-Roth, D. Lenat, and D. Waterman, Rand Corporation, Santa Monica, Ca., 1982.

10. vanMelle, W., Scott, A.C., Bennett, J.S. and Peairs, M.A.S. The EMYCIN Manual. Memo HPP-81-16, Heuristic Programming Project, Stanford University, December 1981.

11. Shortliffe, E.H. ONCOCIN: An aid for the outpatient management of cancer patients. Proceedings of the Society for Computer Medicine, pp 57-58, Washington, D.C., October 1981.